arXiv:1608.05521v1 [cs.PL] 19 Aug 2016

Towards Reversible Computation in Erlang*

Naoki Nishida[1], Adrián Palacios[2,**], and Germán Vidal[2]

[1] Graduate School of Information Science, Nagoya University
Furo-cho, Chikusa-ku, 4648603 Nagoya, Japan,
nishida@is.nagoya-u.ac.jp
[2] MiST, DSIC, Universitat Politècnica de València
Camino de Vera, s/n, 46022 Valencia, Spain
{apalacios, gvidal}@dsic.upv.es

Abstract. In a reversible language, any forward computation can be undone by a finite sequence of backward steps. Reversible computing has been studied in the context of different programming languages and formalisms, where it has been used for debugging and for enforcing fault-tolerance, among others. In this paper, we consider a subset of Erlang, a concurrent language based on the actor model. We formally introduce a reversible semantics for this language. To the best of our knowledge, this is the first attempt to define a reversible semantics for Erlang.

1 Introduction

Let us consider that the operational semantics of a programming language is specified by a state transition relation R such that $R(s, s')$ holds if the state s' is reachable—in one step—from state s. As it is common practice, we let R^* denote the reflexive and transitive closure of R. Then, we say that a programming language (or formalism) is *reversible* if there exists a constructive algorithm that can be used to, given a computation from state s to state s', in symbols $R^*(s, s')$, obtain the state s from s'. In general, such a property does not hold for most programming languages and formalisms. We refer the interested reader to, e.g., [3, 10, 23, 24] for a high level account of the principles of reversible computation.

The notion of *reversible computation* was first introduced in Landauer's seminal work [13] and, then, further improved by Bennett [2] in order to avoid the generation of "garbage" data. The idea underlying these works is that any programming language or formalism can be made reversible by adding the *history* of the computation to each state, which is usually called a *Landauer's embedding*.

* This work has been partially supported by the EU (FEDER) and the Spanish *Ministerio de Economía y Competitividad* (MINECO) under grant TIN2013-44742-C4-1-R and by the *Generalitat Valenciana* under grant PROMETEO-II/2015/013 (SmartLogic). The second and third authors also acknowledge a partial support of COST Action IC1405 on Reversible Computation - extending horizons of computing.

** Partially supported by the the EU (FEDER) and the Spanish *Ayudas para contratos predoctorales para la formación de doctores de la Secretaría de Estado de Investigación, Desarrollo e Innovación del* MINECO under FPI grant BES-2014-069749.

Although carrying the history of a computation might seem infeasible because of its size, there are several successful proposals that are based on this idea. In particular, one can restrict the original language or apply a number of analysis in order to restrict the required information in the history as much as possible, as in, e.g., [17, 18, 21] in the context of a functional language.

Alternatively, one can consider a restricted language where computations are reversible without adding any additional information to the states. This is the case, e.g., of the functional language considered in [25] or the language Janus (see [23] and references therein). Obviously, such languages are not universal since there are functions that cannot be represented (e.g., non-injective functions).

In this paper, we consider the former approach, the so-called Landauer's embedding. In particular, we aim at introducing a form of reversibility in the context of a programming language that follows the actor model (concurrency based on message passing), and that can be considered as a subset of the concurrent and functional language Erlang [1]. Previous approaches have mainly considered reversibility in—mostly synchronous—concurrent calculi like CCS [7, 8], a general framework for reversibility of algebraic process calculi [19], or the recent approach to reversible *session-based* π-calculus [22]. However, we can only find a few approaches that considered the reversibility of *asynchronous* calculi, e.g., Cardelli and Laneve's reversible structures [5] or the approach based on a rollback construct of [11, 14–16] for a higher-order asynchronous π-calculus.

To the best of our knowledge, our work is the first one that considers reversibility in the context of the functional, concurrent, and distributed language Erlang. Here, given a running Erlang system consisting of a pool of interacting processes, possibly distributed in several computers, we aim at allowing a *single* process to undo its actions, including the interactions with other processes, following a rollback fashion. In this context, we must ensure *causal consistency* [7], i.e., an action cannot be undone until all the actions that depend on it have been already undone. E.g., if a process spawns another process, we cannot undo this process spawning until all the actions performed by the new process are undone too. This is particularly challenging in our asynchronous and distributed setting since there is no *global* order for the language events. In this paper, we introduce a rollback operator that undoes the actions of a process until a *checkpoint* is reached. This could be considered as a promising basis for defining *safe* sessions in a language like Erlang.

In this paper, we consider a simple Erlang-like language that can be considered a subset of *Core Erlang* [6]. We present the following contributions: First, we introduce an appropriate standard semantics for the language. In contrast to monolothic previous semantics like that in [4], our semantics is more modular, which simplifies the definition of a reversible extension. In contrast to [20], although we follow some of the ideas in this approach (e.g., the use of a global mailbox), we include the evaluation of expressions and, moreover, our treatment of messages is more deterministic.[3] We then introduce a reversible extension

[3] E.g., in the semantics of [20], at the expression level, the transition semantics of an expression containing a receive statement is, in principle, infinitely branching,

$$
\begin{array}{rcl}
Module & ::= & \mathsf{module}\ Atom = fun_1, \ldots, fun_n \\
fun & ::= & fname = \mathsf{fun}\ (X_1, \ldots, X_n) \rightarrow expr \\
fname & ::= & Atom/Integer \\
lit & ::= & Atom \mid Integer \mid Float \mid [\,] \\
expr & ::= & Var \mid lit \mid fname \mid [expr_1|expr_2] \mid \{expr_1, \ldots, expr_n\} \\
& \mid & \mathsf{call}\ expr\ (expr_1, \ldots, expr_n) \mid \mathsf{apply}\ expr\ (expr_1, \ldots, expr_n) \\
& \mid & \mathsf{case}\ expr\ \mathsf{of}\ clause_1; \ldots; clause_m\ \mathsf{end} \\
& \mid & \mathsf{let}\ Var = expr_1\ \mathsf{in}\ expr_2 \mid \mathsf{receive}\ clause_1; \ldots; clause_n\ \mathsf{end} \\
& \mid & \mathsf{spawn}(expr, [expr_1, \ldots, expr_n]) \mid expr_1\,!\,expr_2 \mid \mathsf{self}() \\
clause & ::= & pat\ \mathsf{when}\ expr_1 \rightarrow expr_2 \\
pat & ::= & Var \mid lit \mid [pat_1|pat_2] \mid \{pat_1, \ldots, pat_n\}
\end{array}
$$

Fig. 1. Language syntax rules

of the standard semantics (basically, a Landauer's embedding). Here, we focus only on the concurrent actions (namely, process spawning, message sending and receiving) and, thus, do not consider the reversibilization of the functional component of the language; rather, we assume that the state of the process—the current expression and its environment—is stored in the history after each execution step. This approach could be improved following, e.g., the approaches from [17, 18, 21]. Finally, we introduce a backward semantics that can be used to undo the actions of a given process in a rollback fashion, until a checkpoint—introduced by the programmer—is reached. Here, ensuring causal consistency is essential and might propagate the rollback action to other, dependent processes.

2 Language Syntax

In this section, we present the syntax of a first-order concurrent and distributed functional language that follows the actor model. Our language is basically equivalent to a subset of Core Erlang [6].

The syntax of the language can be found in Figure 1. Here, a module is a sequence of function definitions, where each function name f/n (atom/arity) has an associated definition of the form $\mathsf{fun}\ (X_1, \ldots, X_n) \rightarrow e$. We consider that a program consists of a single module for simplicity. The body of a function is an *expression*, which can include variables, literals, function names, lists, tuples, calls to built-in functions—mainly arithmetic and relational operators—, function applications, case expressions, let bindings, and receive expressions; furthermore, we also consider the functions spawn, "!" (for sending a message), and self() that are usually considered built-in's in the Erlang language.

Despite the general syntax in Figure 1, as mentioned before, we only consider first order expressions. Therefore, the first expression in calls, applications and spawns can only be function names (instead of arbitrary expressions or closures).

In this language, we distinguish expressions, patterns, and values. As mentioned before, expressions can include all constructs of the language. In contrast,

since their formulation allows for an infinite number of possible queues and selected messages.

patterns are built from variables, literals, lists, and tuples. Finally, *values* are built from literals, lists, and tuples, i.e., they are *ground*—without variables—patterns. Expressions are denoted by $e, e', e_1, e_2, \ldots$, patterns by $p, p', p_1, p_2, \ldots$ and values by $v, v', v_1, v_2, \ldots$ As it is common practice, a *substitution* θ is a mapping from variables to expressions such that $\mathcal{D}om(\theta) = \{X \in \mathit{Var} \mid X \neq \theta(X)\}$ is its domain. Substitutions are usually denoted by sets of mappings like, e.g., $\{X_1 \mapsto v_1, \ldots, X_n \mapsto v_n\}$. Substitutions are extended to morphisms from expressions to expressions in the natural way. The identity substitution is denoted by id. Composition of substitutions is denoted by juxtaposition, i.e., $\theta\theta'$ denotes a substitution θ'' such that $\theta''(X) = \theta'(\theta(X))$ for all $X \in \mathit{Var}$. Also, we denote by $\theta[X_1 \mapsto v_1, \ldots, X_n \mapsto v_n]$ the *update* of θ with the mapping $X_1 \mapsto v_1, \ldots, X_n \mapsto v_n$, i.e., it denotes a new substitution θ' such that $\theta'(X) = v_i$ if $X = X_i$, for some $i \in \{1, \ldots, n\}$, and $\theta'(X) = \theta(X)$ otherwise.

In a case expression "$\mathsf{case}\ e\ \mathsf{of}\ p_1\ \mathsf{when}\ e_1 \to e'_1;\ \ldots;\ p_n\ \mathsf{when}\ e_n \to e'_n\ \ \mathsf{end}$"; we first evaluate e to a value, say v; then, we should find (if any) the first clause $p_i\ \mathsf{when}\ e_i \to e'_i$ such that v matches p_i (i.e., there exists a substitution σ for the variables of p_i such that $v = p_i\sigma$) and $e_i\sigma$—the *guard*—reduces to *true*; then, the case expression reduces to $e'_i\sigma$. Note that guards can only contain calls to built-in functions (typically, arithmetic and relational operators).

As for the concurrent features of the language, we consider that a *system* is of a pool of processes that can only interact through message sending and receiving (i.e., there is no shared memory). Each process has an associated *pid* (process identifier), which is unique in a system. For clarity, we often denote pids with roman letters $\mathrm{p}, \mathrm{p}', \mathrm{p}_1, \ldots$, though they are considered values in our language (i.e., atoms). By abuse of notation, when no confusion can arise, we refer to a process with its pid.

An expression of the form $\mathsf{spawn}(f/n, [e_1, \ldots, e_n])$ has, as a *side effect*, the creation of a new process with a fresh pid p which is initialized with the expression $\mathsf{apply}\ f/n\ (e_1, \ldots, e_n)$; the expression $\mathsf{spawn}(f/n, [e_1, \ldots, e_n])$ itself evaluates to the new pid p. The function $\mathsf{self}()$ just returns the pid of the current process. An expression of the form $\mathrm{p}\,!\,v$ evaluates to the value v and, as a side effect, stores the value v—the *message*—in the queue or *mailbox* of process p.

Finally, an expression "$\mathsf{receive}\ p_1\ \mathsf{when}\ e_1 \to e'_1; \ldots; p_n\ \mathsf{when}\ e_n \to e'_n\ \ \mathsf{end}$" traverses the messages in the process' queue until one of them matches a branch in the receive statement; i.e., it should find the *first* message v in the process' queue (if any) such that $\mathsf{case}\ v\ \mathsf{of}\ p_1\ \mathsf{when}\ e_1 \to e'_1; \ldots; p_n\ \mathsf{when}\ e_n \to e'_n\ \mathsf{end}$ can be reduced; then, the receive expression evaluates to the same expression to which the above case expression would be evaluated, with the additional side effect of deleting the message v from the process' queue. If there is no matching message in the queue, the process suspends its execution until a matching message arrives.

Example 1. Consider the program shown in Figure 2, where the symbol "_" is used to denote an *anonymous* variable, i.e., a variable whose name is not relevant. The computation starts with "$\mathsf{apply}\ \mathit{main}/0\ ()$." Then, this process, say $P1$, spawns two new processes, say $P2$ and $P3$, and then sends the message "*world*"

```
main/0 = fun () → let P2 = spawn(echo/0, [])
                  in let P3 = spawn(target/0, [])
                  in let _ = P3 ! world
                  in let P2 ! {P3, hello}

target/0 = fun () → receive
                      A → receive
                            B → {A, B}
                          end
                    end

echo/0 = fun () → receive
                    {P, M} → P ! M
                  end
```

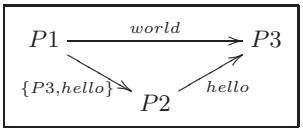

Fig. 2. A simple concurrent program

to process $P3$ and the message $\{P3, hello\}$ to process $P2$, which then resends "$hello$" to $P3$. In our language, there is no guarantee regarding which message arrives first to $P3$, i.e., "apply $main/0$ ()" can evaluate nondeterministically to either $\{hello, world\}$ or $\{world, hello\}$. This is coherent with the semantics of Erlang, where the only guarantee is that if two messages are sent from process p to process p′ and both are delivered, then the order of these messages is kept.

3 The Language Semantics

In order to set precisely the framework for our proposal, in this section we formalize the semantics of the considered language.

Definition 1 (process). *A process is denoted by a tuple of the form* $\langle \mathrm{p}, (\theta, e), q\rangle$ *where* p *is the pid of the process,* (θ, e) *is the control of the state—which consists of an environment (a substitution) and an expression to be evaluated—, and* q *is the process' mailbox, a FIFO queue with the sequence of messages that have been sent to the process.*

A running *system* is then a pool of processes, which is formally defined as follows:

Definition 2 (system). *A system is denoted by* $\Gamma; \Pi$*, where* Γ *is a global mailbox of the system (see below) and* Π *is a pool of processes, denoted by an expression of the form* $\langle \mathrm{p}_1, (\theta_1, e_1), q_1\rangle \,\&\, \cdots \,\&\, \langle \mathrm{p}_\mathrm{n}, (\theta_n, e_n), q_n\rangle$*, where "&" is an associative and commutative operator. We typically denote a system by an expression of the form* $\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi$ *to point out that* $\langle \mathrm{p}, (\theta, e), q\rangle$ *is an arbitrary process of the pool (thanks to the fact that "&" is associative and commutative).*

The role of Γ (which is similar to the "ether" in [20]) will be clarified later, but it is essential to guarantee that all admissible interleavings can be modelled with the semantics. Here, we define Γ as a set of FIFO queues among all (non-necessarily different) pids, i.e., Γ is made of elements of the form $(\mathrm{p}, \mathrm{q}, [v_1, \ldots, v_n])$, where p, q are (not necessarily different) pids and $[v_1, \ldots, v_n]$ is a (possibly empty) ordered list of messages such that v_1 is the oldest message and v_n is the most recent one. For simplicity, we assume that Γ is initialized as follows: $\{(\mathrm{p}, \mathrm{q}, [\,]) \mid \mathrm{p}, \mathrm{q} \text{ are pids}\}$. Then, we use the following notation:

$$(Var)\ \frac{}{\theta, X \xrightarrow{\tau} \theta, \theta(X)} \quad (Tuple)\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e_i' \quad i \in \{1,\ldots,n\}}{\theta, \{e_1,\ldots,e_n\} \xrightarrow{\ell} \theta', \{\overline{e_{1,i-1}}, e_i', \overline{e_{i+1,n}}\}}$$

$$(List1)\ \frac{\theta, e_1 \xrightarrow{\ell} \theta', e_1'}{\theta, [e_1|e_2] \xrightarrow{\ell} \theta', [e_1'|e_2]} \quad (List2)\ \frac{\theta, e_2 \xrightarrow{\ell} \theta', e_2'}{\theta, [e_1|e_2] \xrightarrow{\ell} \theta', [e_1|e_2']}$$

$$(Let1)\ \frac{\theta, e_1 \xrightarrow{\ell} \theta', e_1'}{\theta, \mathsf{let}\ X = e_1\ \mathsf{in}\ e_2 \xrightarrow{\ell} \theta', \mathsf{let}\ X = e_1'\ \mathsf{in}\ e_2} \quad (Let2)\ \frac{}{\theta, \mathsf{let}\ X = v\ \mathsf{in}\ e \xrightarrow{\tau} \theta[X \mapsto v], e}$$

$$(Case1)\ \frac{\theta, e \xrightarrow{\ell} \theta', e'}{\theta, \mathsf{case}\ e\ \mathsf{of}\ cl_1;\ldots;cl_n\ \mathsf{end} \xrightarrow{\ell} \theta', \mathsf{case}\ e'\ \mathsf{of}\ cl_1;\ldots;cl_n\ \mathsf{end}} \quad (Case2)\ \frac{\mathsf{match}(v, cl_1, \ldots, cl_n) = \langle \theta_i, e_i \rangle}{\theta, \mathsf{case}\ v\ \mathsf{of}\ cl_1;\ldots;cl_n\ \mathsf{end} \xrightarrow{\tau} \theta\theta_i, e_i}$$

$$(Call1)\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e_i' \quad i \in \{1,\ldots,n\}}{\theta, \mathsf{call}\ op\ (\overline{e_n}) \xrightarrow{\ell} \theta', \mathsf{call}\ op\ (\overline{e_{1,i-1}}, e_i', \overline{e_{i+1,n}})} \quad (Call2)\ \frac{\mathsf{eval}(op, v_1, \ldots, v_n) = v}{\theta, \mathsf{call}\ op\ (v_1,\ldots,v_n) \xrightarrow{\tau} \theta, v}$$

$$(Apply1)\ \frac{\theta, e_i \xrightarrow{\ell} \theta', e_i' \quad i \in \{1,\ldots,n\}}{\theta, \mathsf{apply}\ a/n\ (\overline{e_n}) \xrightarrow{\ell} \theta', \mathsf{apply}\ a/n\ (\overline{e_{1,i-1}}, e_i', \overline{e_{i+1,n}})}$$

$$(Apply2)\ \frac{\mu(a/n) = \mathsf{fun}\ (X_1,\ldots,X_n) \to e}{\theta, \mathsf{apply}\ a/n\ (v_1,\ldots,v_n) \xrightarrow{\tau} \{X_1 \mapsto v_1, \ldots, X_n \mapsto v_n\}, e}$$

Fig. 3. Standard semantics: evaluation of sequential expressions

$\Gamma \cup (\mathrm{p}, \mathrm{q}, v)$ denotes $\Gamma \setminus \{(\mathrm{p}, \mathrm{q}, vs)\} \cup \{(\mathrm{p}, \mathrm{q}, vs{+}{+}[v]\}$, while $\Gamma \setminus (\mathrm{p}, \mathrm{q}, v)$ denotes $\Gamma \setminus \{(\mathrm{p}, \mathrm{q}, v : vs)\} \cup \{(\mathrm{p}, \mathrm{q}, vs)\}$, where $+{+}$ is the list concatenation operator.

In the following, we denote by $\overline{o_n}$ the sequence of syntactic objects $o_1, \ldots, o_n$ for some n. We also write $\overline{o_{i,j}}$ for the sequence $o_i, \ldots, o_j$ when $i \leq j$ (and the empty sequence otherwise). We write $\overline{o}$ when the number of elements is not relevant.

The semantics is defined by means of two transition relations: $\longrightarrow$ for expressions and $\longmapsto$ for systems. Let us first consider the labelled transition relation

$$\longrightarrow : (Env, Exp) \times Label \times (Env, Exp)$$

where Env and Exp are the domains of environments (i.e., substitutions) and expressions, respectively, and $Label$ denotes an element of the set

$$\{\tau, \mathsf{send}(v_1, v_2), \mathsf{rec}(y, \overline{cl_n}), \mathsf{spawn}(y, a/n, [\overline{e_n}]), \mathsf{self}(y)\}$$

whose meaning will be explained below. For clarity, we divide the transition rules of the semantics for expressions in two sets, depicted in Figures 3 and 4 for sequential and concurrent expressions, respectively.

Most of the rules are self-explanatory. In the following, we only discuss some subtle or complex issues. In principle, the transitions are labelled either with τ (a sequential expression) or with a label that identifies a concurrent action. Labels are used in the system rules (Figure 5) to perform the associated side effects.

$$(Send1)\ \frac{\theta, e_1 \xrightarrow{\ell} \theta', e_1'}{\theta, e_1 \,!\, e_2 \xrightarrow{\ell} \theta', e_1' \,!\, e_2} \qquad \frac{\theta, e_2 \xrightarrow{\ell} \theta', e_2'}{\theta, e_1 \,!\, e_2 \xrightarrow{\ell} \theta', e_1 \,!\, e_2'}$$

$$(Send2)\ \frac{}{\theta, v_1 \,!\, v_2 \xrightarrow{\mathsf{send}(v_1, v_2)} \theta, v_2}$$

$$(Receive)\ \frac{}{\theta, \mathsf{receive}\ cl_1; \ldots; cl_n\ \mathsf{end} \xrightarrow{\mathsf{rec}(y, \overline{cl_n})} \theta, y}$$

$$(Spawn)\ \frac{}{\theta, \mathsf{spawn}(a/n, [e_1, \ldots, e_n]) \xrightarrow{\mathsf{spawn}(y, a/n, [\overline{e_n}])} \theta, y}$$

$$(Self)\ \frac{}{\theta, \mathsf{self}() \xrightarrow{\mathsf{self}(y)} \theta, y}$$

Fig. 4. Standard semantics: evaluation of concurrent expressions

In some of the rules (e.g., for evaluating tuples, lists, etc) we consider for simplicity that elements are evaluated in a non-deterministic way. In an actual programming language the order of evaluation of the arguments in the tuple or list is usually fixed. E.g., in Erlang, reduction takes place from left to right.

For case evaluation, we assume an auxiliary function match which selects the first clause, $cl_i = (p_i\ \mathsf{when}\ e_i' \to e_i)$, such that v matches p_i, i.e., $v = \theta_i(p_i)$, and the guard holds, i.e., $\theta\theta_i, e_i' \longrightarrow^* \theta', true$. Note that, for simplicity, we do not consider here the case in which the argument v matches no clause.

Function calls can either be defined in the program (apply) or be a built-in (call). In the latter case, they are evaluated using the auxiliary function eval. In rule *Apply2*, we consider that the mapping μ stores all function definitions, i.e., it maps every function name a/n to its definition $\mathsf{fun}\ (X_1, \ldots, X_n) \to e$ in the program. As for the applications, note that we only consider first-order functions. In order to extend our semantics to also consider higher-order functions, one should reduce the function name to a *closure* of the form $(\theta', \mathsf{fun}\ (X_1, \ldots, X_n) \to e)$ and, then, reduce e in the environment $\theta'[X_1 \mapsto v_1, \ldots, X_n \mapsto v_n]$. We skip this extension since it is orthogonal to our approach.

Let us now consider the evaluation of concurrent expressions that produce some side effect (Figure 4). Here, we can distinguish two kinds of rules. On the one hand, we have the rules for "!", *Send1* and *Send2*. In this case, we know *locally* what the expression should be reduced to (i.e., v_2 in rule *Send2*). For the remaining rules, this is not known locally and, thus, we return a fresh distinguished symbol, $y \notin Var$ (by abuse, y is dealt with as a variable), so that the system rules will eventually bind y to its correct value. This *trick* allows us to keep the rules for expressions and systems separated (i.e., the semantics shown in Figures 3 and 4 is mostly independent from the rules in Figure 5), in contrast to other calculi, e.g., [4], where they are combined into a single transition relation.

Let us finally consider the system rules, which are depicted in Figure 5. In most of the rules, we consider an arbitrary system of the form $\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi$, where Γ is the global mailbox and $\langle \mathrm{p}, (\theta, e), q\rangle \& \Pi$ is a pool of process that contains at least one process $\langle \mathrm{p}, (\theta, e), q\rangle$.

$$(Exp)\qquad \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi \longmapsto \Gamma; \langle \mathrm{p}, (\theta', e'), q\rangle \& \Pi}$$

$$(Send)\qquad \frac{\theta, e \xrightarrow{\mathsf{send}(\mathrm{p}'', v)} \theta', e'}{\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi \longmapsto \Gamma \cup (\mathrm{p}, \mathrm{p}'', v); \langle \mathrm{p}, (\theta', e'), q\rangle \& \Pi}$$

$$(Receive)\qquad \frac{\theta, e \xrightarrow{\mathsf{rec}(y, \overline{cl_n})} \theta', e' \quad \mathsf{matchrec}(\overline{cl_n}, q) = (\theta_i, e_i, q')}{\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi \longmapsto \Gamma; \langle \mathrm{p}, (\theta'\theta_i, e'\{y \mapsto e_i\}), q'\rangle \& \Pi}$$

$$(Spawn)\qquad \frac{\theta, e \xrightarrow{\mathsf{spawn}(y, a/n, [\overline{e_n}])} \theta', e' \quad \mathrm{p}' \text{ is a fresh pid}}{\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi \longmapsto \Gamma; \langle \mathrm{p}, (\theta', e'\{y \mapsto \mathrm{p}'\}), q\rangle \& \langle \mathrm{p}', (\theta', \mathsf{apply}\ a/n\ (e_1, \ldots, e_n)), [\,]\rangle \& \Pi}$$

$$(Self)\qquad \frac{\theta, e \xrightarrow{\mathsf{self}(y)} \theta', e'}{\Gamma; \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi \longmapsto \Gamma; \langle \mathrm{p}, (\theta', e'\{y \mapsto \mathrm{p}\}), q\rangle \& \Pi}$$

$$(Sched)\qquad \frac{\alpha(\Gamma) = (\mathrm{p}', \mathrm{p}) \quad \Pi = \langle \mathrm{p}, (\theta, e), q\rangle \& \Pi'}{\Gamma; \Pi \longmapsto \Gamma \setminus \{(\mathrm{p}', \mathrm{p}, v)\}; \langle \mathrm{p}, (\theta, e), v : q\rangle \& \Pi'}$$

Fig. 5. Standard semantics: system rules

Note that, in rule *Send*, we add the triple $(\mathrm{p}, \mathrm{p}'', v)$ to Γ instead of adding it to the queue of process p''. This is necessary to ensure that all possible non-deterministic results can be obtained (as discussed in Example 1). Observe that e' is usually different from v since e may have different nested operators. E.g., if e has the form "$\mathsf{case}\ \mathrm{p}\,!\,v\ \mathsf{of}\ \{\ldots\}$," then e' will be "$\mathsf{case}\ v\ \mathsf{of}\ \{\ldots\}$" with label $\mathsf{send}(\mathrm{p}, v)$.

In rule *Receive*, we use the auxiliary function matchrec to evaluate a receive expression. The main difference with match is that matchrec also takes a queue q and returns the modified queue q'. Then, the distinguished variable y is bound to the expression in the selected clause, e_i, and the environment is extended with the matching substitution. If no message in the queue q matches any clause, then the rule is not applicable and the selected process cannot be reduced (i.e., it suspends).

With the rules presented so far, any system will soon reach a state in which no reduction can be performed, since messages are stored in the global mailbox, but they are not dispatched to the queues of the processes. This is precisely the task of the scheduler, which is modelled by rule *Sched*. The rule is non-deterministic, so any scheduling policy can be modelled by the semantics. A message is selected from the list of messages by the auxiliary function α, which can select any arbitrary pair of (non-necessarily different) pids $(\mathrm{p}', \mathrm{p})$. Note that we take the oldest message in the queue—the first one in the list—, which is necessary to ensure that "the messages sent—directly—between two given processes arrive in the same order they were sent", as mentioned in the previous section.

Example 2. Let us consider the program shown in Figure 6 and a possible execution trace. This trace is modelled by our semantics. For clarity, we only show

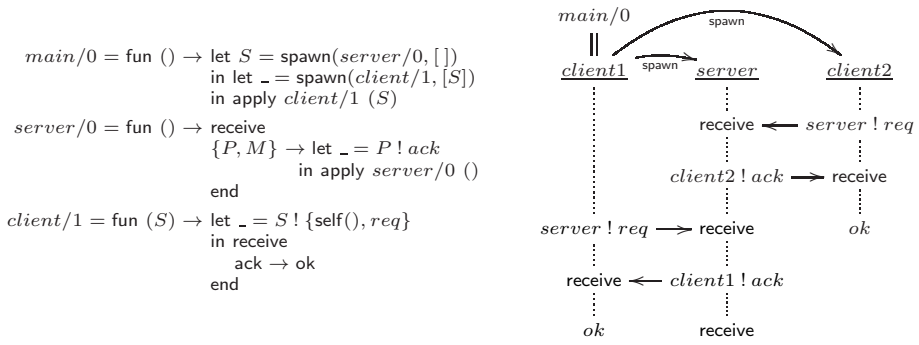

Fig. 6. A simple client-server

$$
\begin{aligned}
& [\,];\ \langle \mathrm{c1}, (id, \underline{\mathsf{apply}\ main/0\ ()}), [\,]\rangle \\
\longmapsto\ & \ldots \\
\longmapsto\ & [\,];\ \langle \mathrm{c1}, (\sigma, \mathsf{let}\ _ = \underline{S\,!\,\{\mathrm{c1}, req\}}\ \mathsf{in}\ \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [(\mathrm{c1}, \mathrm{s}, [v_2])];\ \langle \mathrm{c1}, (\sigma, \underline{\mathsf{let}\ _ = \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots}), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & \underline{[(\mathrm{c1}, \mathrm{s}, [v_2])}];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [\,];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \underline{\mathsf{receive}\ \{P, M\} \to \ldots}), [\{\mathrm{c1}, req\}]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [\,];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (\theta_2, \mathsf{let}\ _ = \underline{P\,!\,ack}\ \mathsf{in}\ \ldots), [\{\mathrm{c1}, req\}]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [(\mathrm{s}, \mathrm{c1}, [ack])];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (\theta_2, \underline{\mathsf{let}\ _ = ack\ \mathsf{in}\ \ldots}), [\{\mathrm{c1}, req\}]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [(\mathrm{s}, \mathrm{c1}, [ack])];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (\theta_2, \underline{\mathsf{apply}\ server/0\ ()}), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & \underline{[(\mathrm{s}, \mathrm{c1}, [ack])}];\ \langle \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [\,];\ \langle \mathrm{c1}, (\sigma, \underline{\mathsf{receive}\ ack \to \ldots}), [ack]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\longmapsto\ & [\,];\ \langle \mathrm{c1}, (\sigma, ok), [\,]\rangle\ \&\ \langle \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
& \qquad \langle \mathrm{c2}, (\sigma, ok), [\,]\rangle
\end{aligned}
$$

Fig. 7. A trace with $\sigma = \{S \mapsto \mathrm{s}\}$, $\theta_2 = \{P \mapsto \mathrm{c1}, M \mapsto req\}$, and $v_2 = \{\mathrm{c1}, req\}$.

in Figure 7 the transition steps that correspond to the last two messages sent between *client1* and *server*.

4 Reversible Semantics

In this section, we introduce a reversible extension of the semantics defined so far. Moreover, thanks to our modular design, the semantics for the language expressions needs not be changed.

To be precise, in this section we introduce two transition relations: $\rightharpoonup$ and $\leftharpoondown$. The first relation, $\rightharpoonup$, is a conservative extension of the standard semantics $\longmapsto$ (Figure 5) to also include some additional information in the states, following a typical Landauer's embedding. We refer to $\rightharpoonup$ as the *forward* reversible semantics (or simply the forward semantics). In contrast, the second relation, $\leftharpoondown$, proceeds in the backwards direction, "undoing" the actions of a single process (and, by

causal consistency, possibly propagating it to other processes). We refer to $\leftharpoondown$ as the backward (reversible) semantics. Finally, we denote the union $\rightharpoonup \cup \leftharpoondown$ by $\rightleftharpoons$. Then, in a computation modelled with $\rightleftharpoons$ the system mainly evolves forwards, except for some processes that can run backwards in order to undo some particular actions (and, afterwards, will run forwards again).

Here, we will introduce a (non-deterministic) "undo" operation (cf. rules *Undo1* and *Undo2* in Figure 10) which has some similarities to, e.g., the rollback operator of [11]. In order to delimit the scope of this operation (i.e., to determine when to stop undoing the actions of a process), we allow the programmer to introduce *checkpoints* in a program. Syntactically, they are denoted with the built-in function check, which is used in a let expression as follows: "let _ = check in *expr*". Function check returns an identifier t associated to the checkpoint (see below). In the following, we consider that the rules to evaluate the language expressions (Figures 3 and 4) are extended with the following rule:

$$(\mathit{Check}) \quad \frac{}{\theta, \mathsf{check} \xrightarrow{\mathsf{check(y)}} \theta, \mathsf{y}}$$

In the next section, we will see that the only effect of a call to function check is to add a checkpoint in the trace of a given process.

4.1 Forward Semantics

First, we introduce the forward (reversible) semantics. Since the expression rules are the same (except for the additional rule for check mentioned above), we will only introduce the reversible system rules, which are shown in Figure 8. Processes now include a memory π that records the intermediate states of a process. Note that we could optimize this in order to follow a strategy similar to that in [17, 18, 21] for the reversibility of functional expressions, but this is orthogonal to our purpose in this paper, so we focus only on the concurrent actions.

The rules are mostly self-explanatory. Here, we use a special symbol $\#_k^{\mathtt{t}}$ for checkpoints, where k denotes the type of checkpoint and t is a (unique) identifier for the checkpoint. The checkpoints introduced by the programmer, denoted by $\#_{\mathsf{ch}}^{\mathtt{t}}$, represent a *safe* point in the program execution, i.e., if we are undoing the actions of a given process, we can safely stop the backward computation at any checkpoint. Rollback operations (i.e., rules *Undo1* and *Undo2* in Figure 10) and checkpoints introduced by the programmer lay the ground for defining *safe* sessions whose actions can be undone if, e.g., an exception occurs before they are completed. In practice, we distinguish three types of checkpoints:

- Checkpoints introduced by the programmer. These checkpoints correspond to the evaluation of function check and are denoted with $\#_{\mathsf{ch}}$.
- Checkpoints associated to receiving a message ($\#_\alpha$) and spawning a process ($\#_{\mathsf{sp}}$). These checkpoints are internal and only used to ensure causal consistency.

Also, we now assume that all messages v are transformed into a tuple $\{\mathtt{t}, v\}$ where t is a unique identifier for this message. Consequently, receive expressions

$$(\textit{Internal}) \quad \frac{\theta, e \xrightarrow{\tau} \theta', e'}{\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma; \langle \tau(\theta, e) : \pi, \mathrm{p}, (\theta', e'), q\rangle \& \Pi}$$

$$(\textit{Check}) \quad \frac{\theta, e \xrightarrow{\mathsf{check}(y)} \theta', e' \quad \text{and } \mathtt{t} \text{ is fresh}}{\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma; \langle \mathsf{check}(\theta, e) : \#^{\mathtt{t}}_{\mathsf{ch}} : \pi, \mathrm{p}, (\theta', e'\{y \mapsto \mathtt{t}\}), q\rangle \& \Pi}$$

$$(\textit{Receive}) \quad \frac{\theta, e \xrightarrow{\mathsf{rec}(y, \overline{cl_n})} \theta', e' \quad \mathsf{matchrec}(\overline{cl_n}, q) = (\theta_i, e_i, q', \mathtt{t})}{\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma; \langle \mathsf{rec}(\theta, e, q) : \pi, \mathrm{p}, (\theta'\theta_i, e'\{y \mapsto e_i\}), q'\rangle \& \Pi}$$

$$(\textit{Send}) \quad \frac{\theta, e \xrightarrow{\mathsf{send}(\mathrm{p}'', v)} \theta', e' \quad \text{and } \mathtt{t} \text{ is fresh}}{\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma \cup (\mathrm{p}, \mathrm{p}'', \{\mathtt{t}, v\}); \langle \mathsf{send}(\mathrm{p}'', \theta, e, \mathtt{t}) : \pi, \mathrm{p}, (\theta', e'), q\rangle \& \Pi}$$

$$(\textit{Spawn}) \quad \frac{\theta, e \xrightarrow{\mathsf{spawn}(y, a/n, [e_1, \ldots, e_n])} \theta', e' \quad \mathrm{p}' \text{ is a fresh pid} \quad \text{and } \mathtt{t} \text{ is fresh}}{\begin{array}{r}\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma; \langle \mathsf{spawn}(\theta, e, \mathrm{p}') : \pi, \mathrm{p}, (\theta', e'\{y \mapsto \mathrm{p}'\}), q\rangle \\ \& \langle [\,], \mathrm{p}', (\theta, (\mathsf{apply}\ a/n\ (e_1, \ldots, e_n))), [\,]\rangle \& \Pi\end{array}}$$

$$(\textit{Self}) \quad \frac{\theta, e \xrightarrow{\mathsf{self}(y)} \theta', e'}{\Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi \rightharpoonup \Gamma; \langle \mathsf{self}(\theta, e) : \pi, \mathrm{p}, (\theta', e'\{y \mapsto \mathrm{p}\}), q\rangle \& \Pi}$$

$$(\textit{Sched}) \quad \frac{\alpha(\Gamma) = (\mathrm{p}', \mathrm{p}) \qquad \Pi = \langle \pi, \mathrm{p}, (\theta, e), q\rangle \& \Pi'}{\Gamma; \Pi \rightharpoonup \Gamma \setminus (\mathrm{p}', \mathrm{p}, \{\mathtt{t}, v\}); \langle \alpha(\mathrm{p}', \mathrm{p}, \{\mathtt{t}, v\}) : \pi, \mathrm{p}, (\theta, e), \{\mathtt{t}, v\} : q\rangle \& \Pi'}$$

Fig. 8. Reversible semantics: system rules

should ignore this part of the message when matching with the corresponding patterns.

Example 3. Consider again the program shown in Figure 6, where the function $client/1$ is now defined as follows:

$$\begin{array}{l} client/1 = \mathsf{fun}\ (S) \rightarrow \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \mathsf{let}\ _ = S\,!\,\{\mathsf{self}(), req\} \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \mathsf{in}\ \mathsf{receive}\ \mathsf{ack} \rightarrow \mathsf{ok}\ \mathsf{end} \end{array}$$

and the execution trace shown in Figure 7. The corresponding forward (reversible) computation is shown in Figure 9.

4.2 Backward Semantics

Now, we present the backward semantics. The rules can be split into two groups: the ones required to trigger, manage, and finish a rollback operation, and the ones that actually perform the backward computation.

We denote a process running backwards with $\lfloor \mathrm{p} \rfloor_{\Psi}$, where Ψ is the set of *pending* checkpoints that the backward computation of p has to go through before resuming its forward computation. For instance, a process of the form $\lfloor \mathrm{p} \rfloor_{\{\#^{\mathtt{t}}_{\mathsf{ch}}\}}$ should go backwards until a checkpoint $\#^{\mathtt{t}}_{\mathsf{ch}}$ is found in its trace, a process $\lfloor \mathrm{p} \rfloor_{\{\#^{\mathtt{t}}_{\alpha}\}}$ should go backwards until an event of the form $\alpha(\ldots, \{\mathtt{t}, _\})$ is found in its trace, and a process $\lfloor \mathrm{p} \rfloor_{\{\#^{\mathtt{t}}_{\mathsf{sp}}\}}$ should go backwards until its initial

$$
\begin{array}{ll}
 & [\,];\ \langle [\,], \mathrm{c1}, (id, \underline{\mathsf{apply}\ main/0\ ()}), [\,]\rangle \\
\rightharpoonup & \ldots \\
\rightharpoonup & [\,];\ \langle \pi_i, \mathrm{c1}, (\sigma, \mathsf{let}\ _ = \underline{\mathsf{check}}\ \mathsf{in}\ \ldots), [\,]\rangle\ \&\ \langle \pi_i', \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \& \\
 & \quad \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [\,];\ \langle \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \mathrm{c1}, (\sigma, \underline{\mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots}), [\,]\rangle\ \& \\
 & \quad \langle \pi_i', \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \&\ \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [\,];\ \langle \tau(\sigma, \mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots) : \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \\
 & \quad \mathrm{c1}, (\sigma, \mathsf{let}\ _ = \underline{S\ !\ \{\mathrm{c1}, req\}}\ \mathsf{in}\ \ldots), [\,]\rangle\ \& \\
 & \quad \langle \pi_i', \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \&\ \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [(\mathrm{c1}, \mathrm{s}, [\{\mathtt{t}_2, v_2\}])];\ \langle \mathsf{send}(\mathrm{s}, \sigma, \mathsf{let}\ _ = S\ !\ \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots, \mathtt{t}_2) : \tau(\sigma, \mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots) \\
 & \quad : \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \mathrm{c1}, (\sigma, \underline{\mathsf{let}\ _ = \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots}), [\,]\rangle\ \& \\
 & \quad \langle \pi_i', \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \&\ \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [\underline{(\mathrm{c1}, \mathrm{s}, [\{\mathtt{t}_2, v_2\}])}];\ \langle \tau(\sigma, \mathsf{let}\ _ = \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots) : \mathsf{send}(\mathrm{s}, \sigma, \mathsf{let}\ _ = S\ !\ \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots, \mathtt{t}_2) : \tau(\sigma, \mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots) \\
 & \quad : \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \& \\
 & \quad \langle \pi_i', \mathrm{s}, (id, \mathsf{receive}\ \{P, M\} \to \ldots), [\,]\rangle\ \&\ \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [\,];\ \langle \tau(\sigma, \mathsf{let}\ _ = \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots) : \mathsf{send}(\mathrm{s}, \sigma, \mathsf{let}\ _ = S\ !\ \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots, \mathtt{t}_2) : \tau(\sigma, \mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots) \\
 & \quad : \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \& \\
 & \quad \langle \alpha(\mathrm{c1}, \mathrm{s}, \{\mathtt{t}_2, v_2\}) : \pi_i', \mathrm{s}, (id, \underline{\mathsf{receive}\ \{P, M\} \to \ldots}), [\{\mathtt{t}_2, \{\mathrm{c1}, req\}\}]\rangle\ \& \\
 & \quad \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & [\,];\ \langle \tau(\sigma, \mathsf{let}\ _ = \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots) : \mathsf{send}(\mathrm{s}, \sigma, \mathsf{let}\ _ = S\ !\ \{\mathrm{c1}, req\}\ \mathsf{in}\ \ldots, \mathtt{t}_2) : \tau(\sigma, \mathsf{let}\ _ = \mathtt{t}_1\ \mathsf{in}\ \ldots) \\
 & \quad : \mathsf{check}(\sigma, \mathsf{let}\ _ = \mathsf{check}\ \mathsf{in}\ \ldots) : \#_{\mathsf{ch}}^{\mathtt{t}_1} : \pi_i, \mathrm{c1}, (\sigma, \mathsf{receive}\ ack \to \ldots), [\,]\rangle\ \& \\
 & \quad \langle \mathsf{rec}(id, \mathsf{receive}\ \{P, M\} \to \ldots, [\{\mathtt{t}_2, \{\mathrm{c1}, req\}\}]) : \alpha(\mathrm{c1}, \mathrm{s}, \{\mathtt{t}_2, v_2\}) : \pi_i', \\
 & \quad \mathrm{s}, (\theta_2, \mathsf{let}\ _ = \underline{P\ !\ ack}\ \mathsf{in}\ \ldots), [\,]\rangle\ \&\ \langle \pi_i'', \mathrm{c2}, (\sigma, ok), [\,]\rangle \\
\rightharpoonup & \ldots
\end{array}
$$

Fig. 9. A (partial) trace with the forward reversible semantics.

state is reached (i.e., it should be completely undone). Furthermore, a process may have a set with several pending checkpoints coming from different rollback operations. Once a checkpoint is reached, it is removed from set Ψ. When $\Psi = \emptyset$, the process p can resume its forward computation (from the last checkpoint).

Let us now explain the rules shown in Figure 10, whose purpose is to mark (or unmark) a process for backward computation. The first two rules, *Undo1* and *Undo2*, are used to introduce a new rollback for a checkpoint $\#_{\mathsf{ch}}^{\mathtt{t}}$ for some $\mathtt{t}$.

A backward computation ends when there are no pending checkpoints in the set (rule *Stop1*) or when we reach a process with an empty trace and a spawn checkpoint $\#_{\mathsf{sp}}^{\mathtt{t}}$ in the set of pending checkpoints, so that the process is removed from the system (rule *Stop2*).

Finally, if we reach a checkpoint $\#_{\mathsf{ch}}^{\mathtt{t}}$ in the trace of a process running backwards, we can either remove it from the set of pending checkpoints (rule *Check*) or just ignore it when it was not in the set of pending checkpoints (rule *Discard*).[4]

Let us now discuss the rules for performing backward computations, which are shown in Figure 11. In general, these rules are applied whenever the term located at the top of the computation history π is an event associated to a sequential or concurrent expression (i.e., when it is not a checkpoint). These terms are deleted from the trace when any of these rules is applied.

In general, all rules restore the control (and sometimes also the queue) of a process. In order to reverse the creation of a process, rule *Spawn* marks the child

[4] Note that this may happen often since the evaluation of each function check introduces a checkpoint $\#_{\mathsf{ch}}^{\mathtt{t}}$ no matter whether a rollback operation is run or not.

$$(\mathit{Undo1}) \quad \Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \,\&\, \Pi \leftharpoonup \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\#^{\mathtt{t}}_{\mathsf{ch}}} \,\&\, \Pi$$

$$(\mathit{Undo2}) \quad \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi} \,\&\, \Pi \leftharpoonup \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi \cup \#^{\mathtt{t}}_{\mathsf{ch}}} \,\&\, \Pi$$

$$(\mathit{Stop1}) \quad \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\emptyset} \,\&\, \Pi \leftharpoonup \Gamma; \langle \pi, \mathrm{p}, (\theta, e), q\rangle \,\&\, \Pi$$

$$(\mathit{Stop2}) \quad \Gamma; \lfloor \langle [\,], \mathrm{p}, (\theta, e), q\rangle \rfloor_{\{\#^{\mathtt{t}}_{\mathsf{sp}}\}} \,\&\, \Pi \leftharpoonup \Gamma; \Pi$$

$$(\mathit{Check}) \quad \Gamma; \lfloor \langle \#^{\mathtt{t}}_{\mathsf{ch}} : \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi \cup \#^{\mathtt{t}}_{\mathsf{ch}}} \,\&\, \Pi \leftharpoonup \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi} \,\&\, \Pi$$

$$(\mathit{Discard}) \quad \Gamma; \lfloor \langle \#^{\mathtt{t}}_{\mathsf{ch}} : \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi} \,\&\, \Pi \leftharpoonup \Gamma; \lfloor \langle \pi, \mathrm{p}, (\theta, e), q\rangle \rfloor_{\Psi} \,\&\, \Pi \quad \text{if } \#^{\mathtt{t}}_{\mathsf{ch}} \notin \Psi$$

Fig. 10. Backward semantics: rules for dealing with rollbacks

process with checkpoint $\#^{\mathrm{p}'}_{\mathsf{sp}}$. The child process will then run backwards until, eventually, its trace is empty and rule *Stop2* removes it from the system.

For undoing the sending of a message, rule *Send1* removes a message from Γ when the message has not been delivered yet. Here, we use the operator "\\" to denote the removal of a message no matter its position (in contrast to "\" which always removes the oldest message of the list, i.e., the first one). Otherwise, rule *Send2* propagates the backward computation to the receiver process by adding $\#^{\mathtt{t}}_{\alpha}$ as a pending checkpoint of the receiver process ($\mathtt{t}$ identifies the message). This will cause the receiver process to undo all the actions that it has performed since it received the message, thus ensuring *causal consistency*.

Observe that, at first glance, one may think rule *Receive* should also introduce some new rollback operation for causal consistency. However, if we take a closer look, we will realize that receiving a message in our context is just about processing the message, and not actually receiving it. In fact, the processed message could have been delivered to the process mailbox a long time ago, and triggering a backward computation on the sending process would be unnecessary.

It is actually rules *Sched1* and *Sched2* that take care of dealing with terms of the form $\alpha(\mathrm{p}'', \mathrm{p}, \{\mathtt{t}, v\})$ in a trace. When the rollback was introduced by another process (in rule *Send2*), rule *Sched1* removes the corresponding message $\{\mathtt{t}, v\}$ from the process queue and deletes the checkpoint $\#^{\mathtt{t}}_{\alpha}$ from the set of pending checkpoints. No further action is required; actually, when the process p'' that triggered rule *Send2* resumes its forward computation, the same message v (with a different identifier $\mathtt{t}'$) will be delivered again to process p. When the rollback was not introduced using rule *Send2*, i.e., there is no checkpoint $\#^{\mathtt{t}}_{\alpha}$ in Ψ, then *Sched2* proceeds as before but also adds the message $(\mathrm{p}'', \mathrm{p}, \{\mathtt{t}, v\})$ to the global mailbox Γ. Then, when this process resumes its forward computation, the same message will again be delivered by the scheduler. Here, the operator $\overline{\cup}$ is used to add messages to the head of the corresponding list instead of to its end, i.e., $\Gamma \,\overline{\cup}\, (\mathrm{p}, \mathrm{q}, \{\mathtt{t}, v\})$ denotes $\Gamma \setminus \{(\mathrm{p}, \mathrm{q}, vs)\} \cup \{(\mathrm{p}, \mathrm{q}, \{\mathtt{t}, v\} : vs)\}$.

In both cases, we aim at eventually producing a system that could have been obtained from the initial state using only forward steps.

$$(\mathit{Internal}) \quad \Gamma; \lfloor\langle \tau(\theta, e) : \pi, p, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \Pi \leftharpoonup \Gamma; \lfloor\langle \pi, p, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi$$

$$(\mathit{Check}) \quad \Gamma; \lfloor\langle \mathsf{check}(\theta, e) : \pi, p, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \Pi \leftharpoonup \Gamma; \lfloor\langle \pi, p, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi$$

$$(\mathit{Receive}) \quad \Gamma; \lfloor\langle \mathsf{rec}(\theta, e, q) : \pi, \mathrm{p}, (\theta', e'), q'\rangle\rfloor_\Psi \;\&\; \Pi \leftharpoonup \Gamma; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi$$

$$(\mathit{Send1}) \quad \begin{array}{r} \Gamma; \lfloor\langle \mathsf{send}(\mathrm{p}'', \theta, e, \mathtt{t}) : \pi, \mathrm{p}, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \Pi \leftharpoonup \Gamma'; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi \\ \text{if } (\mathrm{p}, \mathrm{p}'', \{\mathtt{t}, v\}) \text{ occurs in } \Gamma, \text{ with } \Gamma' = \Gamma \setminus\!\setminus (\mathrm{p}, \mathrm{p}'', v, \mathtt{t}) \end{array}$$

$$(\mathit{Send2}) \quad \begin{array}{r} \Gamma; \lfloor\langle \mathsf{send}(\mathrm{p}'', \theta, e, \mathtt{t}) : \pi, \mathrm{p}, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \langle \pi'', \mathrm{p}'', (\theta'', e''), q''\rangle \;\&\; \Pi \\ \leftharpoonup \Gamma; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \lfloor\langle \pi'', \mathrm{p}'', (\theta'', e''), q''\rangle\rfloor_{\#^{\mathtt{t}}_\alpha} \;\&\; \Pi \\ \text{if } (\mathrm{p}, \mathrm{p}'', \{\mathtt{t}, v\}) \text{ does not occur in } \Gamma \end{array}$$

$$(\mathit{Spawn}) \quad \begin{array}{r} \Gamma; \lfloor\langle \mathsf{spawn}(\theta, e, \mathrm{p}'') : \pi, \mathrm{p}, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \langle \pi'', \mathrm{p}'', (\theta'', e''), q''\rangle\rangle \;\&\; \Pi \\ \leftharpoonup \Gamma; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \lfloor\langle \pi'', \mathrm{p}'', (\theta'', e''), q''\rangle\rangle\rfloor_{\#^{\mathrm{p}''}_{\mathsf{sp}}} \;\&\; \Pi \end{array}$$

$$(\mathit{Self}) \quad \Gamma; \lfloor\langle \mathsf{self}(\theta, e) : \pi, \mathrm{p}, (\theta', e'), q\rangle\rfloor_\Psi \;\&\; \Pi \leftharpoonup \Gamma; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi$$

$$(\mathit{Sched1}) \quad \Gamma; \lfloor\langle \alpha(\mathrm{p}'', \mathrm{p}, \{\mathtt{t}, v\}) : \pi, \mathrm{p}, (\theta, e), \{\mathtt{t}, v\} : q\rangle\rfloor_{\Psi \cup \#^{\mathtt{t}}_\alpha} \;\&\; \Pi \leftharpoonup \Gamma; \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi$$

$$(\mathit{Sched2}) \quad \begin{array}{r} \Gamma; \lfloor\langle \alpha(\mathrm{p}'', \mathrm{p}, \{\mathtt{t}, v\}) : \pi, \mathrm{p}, (\theta, e), \{\mathtt{t}, v\} : q\rangle\rfloor_\Psi \;\&\; \Pi \\ \leftharpoonup \Gamma \,\overline{\cup}\, (\mathrm{p}'', \mathrm{p}, \{\mathtt{t}, v\}); \lfloor\langle \pi, \mathrm{p}, (\theta, e), q\rangle\rfloor_\Psi \;\&\; \Pi \quad \text{if } \#^{\mathtt{t}}_\alpha \notin \Psi \end{array}$$

Fig. 11. Backward semantics: Rules for backward computation.

Example 4. Consider the forward execution trace shown in Figure 9. A corresponding backward computation is shown in Figure 12.

The soundness of our reversible semantics can be stated as follows: i) every forward step can be reversed; ii) every system reached during a computation with $\rightleftharpoons$, could have been reached with $\rightharpoonup$ from the initial system.

5 Discussion

We have defined a reversible semantics for a first-order subset of Erlang. To the best of our knowledge, this is the first attempt to define a reversible semantics for Erlang. As mentioned in the introduction, the closest to our work is the debugging approach based on a rollback construct of [11, 14–16], but it is defined in the context of a higher-order asynchronous π-calculus. Also, we share some similarities with the checkpointing technique for fault-tolerant distributed computing of [9, 12], although the aim is different (they aim at defining a new language rather than extending an existing one).

As future work, we plan to formally study the correctness of our calculi, as claimed at the end of the previous section. Also, we would like to extend the Erlang language with a new construct for *safe sessions* so that all the actions in a session can be undone when the session aborts, which has a great potential to automate the fault-tolerance capabilities of the language Erlang.

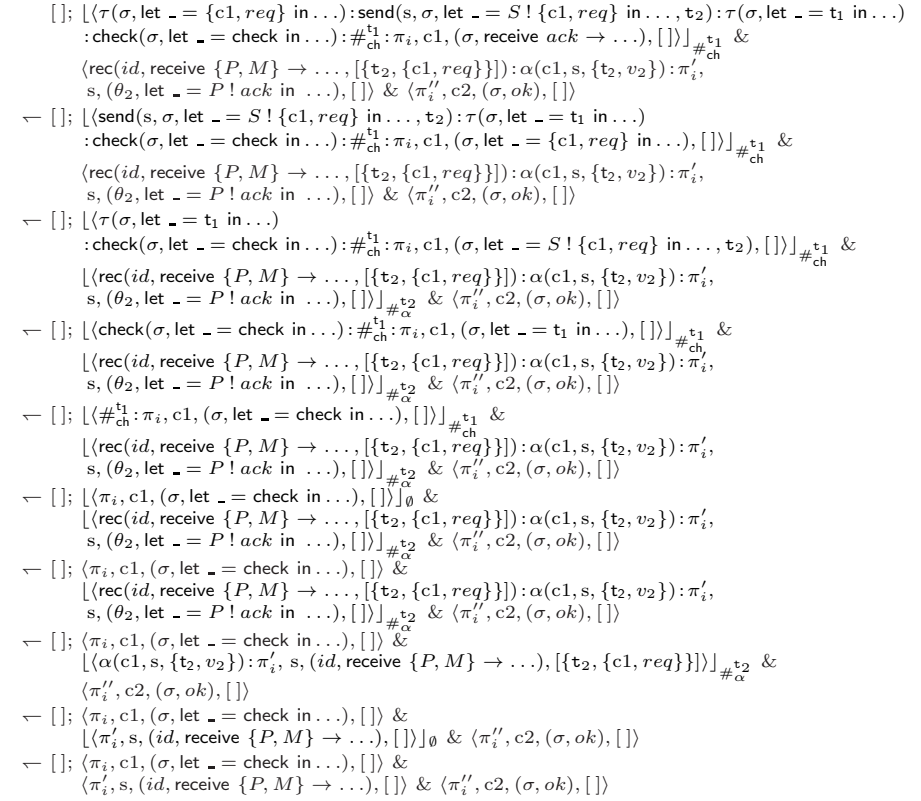

Fig. 12. A (partial) trace with the backward reversible semantics.

Acknowledgements

We would like to thank Ivan Lanese and the anonymous reviewers for many useful suggestions to improve this paper.

References

1. Armstrong, J., Virding, R., Williams, M.: Concurrent programming in Erlang (2nd edition). Prentice Hall (1996)
2. Bennett, C.: Logical reversibility of computation. IBM Journal of Research and Development 17, 525–532 (1973)
3. Bennett, C.: Notes on the history of reversible computation. IBM Journal of Research and Development 44(1), 270–278 (2000)
4. Caballero, R., Martín-Martín, E., Riesco, A., Tamarit, S.: A declarative debugger for concurrent erlang programs (extended version). Tech. Rep. SIC-15/13, UCM (2013), `http://maude.sip.ucm.es/~adrian/files/conc_cal_tr.pdf`
5. Cardelli, L., Laneve, C.: Reversible structures. In Proc. of CMSB 2011, pp. 131–140. ACM (2011)

6. Carlsson, R., Gustavsson, B., Johansson, E. *et al*: Core erlang 1.0.3. language specification (2004), available from `https://www.it.uu.se/research/group/hipe/cerl/doc/core_erlang-1.0.3.pdf`
7. Danos, V., Krivine, J.: Reversible communicating systems. In Proc. of CONCUR 2004. Springer LNCS 3170, pp. 292–307 (2004)
8. Danos, V., Krivine, J.: Transactions in RCCS. In Proc. of CONCUR 2005. Springer LNCS 3653, pp. 398–412 (2005)
9. Field, J., Varela, C.A.: Transactors: a programming model for maintaining globally consistent distributed state in unreliable environments. In Proc. of POPL 2005, pp. 195–208. ACM (2005)
10. Frank, M.P.: Introduction to reversible computing: motivation, progress, and challenges. In Proc. of 2nd Conf. on Computing Frontiers, pp. 385–390. ACM (2005)
11. Giachino, E., Lanese, I., Mezzina, C.A.: Causal-consistent reversible debugging. In Proc. of FASE 2014. Springer LNCS 8411, pp. 370–384 (2014)
12. Kuang, P., Field, J., Varela, C.A.: Fault tolerant distributed computing using asynchronous local checkpointing. In Proc. of AGERE! 2014, pp. 81–93. ACM (2014)
13. Landauer, R.: Irreversibility and heat generation in the computing process. IBM Journal of Research and Development 5, 183–191 (1961)
14. Lanese, I., Mezzina, C.A., Schmitt, A., Stefani, J.: Controlling reversibility in higher-order pi. In Proc. CONCUR 2011. Springer LNCS 6901, pp. 297–311 (2011)
15. Lanese, I., Mezzina, C.A., Stefani, J.: Reversibility in the higher-order π-calculus. Theor. Comput. Sci. 625, 25–84 (2016)
16. Lienhardt, M., Lanese, I., Mezzina, C.A., Stefani, J.: A reversible abstract machine and its space overhead. In Proc. of the Joint FMOODS 2012 and FORTE 2012 Int'l Conference. Springer LNCS 7273, pp. 1–17 (2012)
17. Matsuda, K., Hu, Z., Nakano, K., Hamana, M., Takeichi, M.: Bidirectionalization transformation based on automatic derivation of view complement functions. In Proc. of ICFP 2007, pp. 47–58. ACM (2007)
18. Nishida, N., Palacios, A., Vidal, G.: Reversible term rewriting. In: Kesner, D., Pientka, B. (eds.) Proc. of the First International Conference on Formal Structures for Computation and Deduction (FSCD'16). Leibniz International Proceedings in Informatics (2016)
19. Phillips, I., Ulidowski, I.: Reversing algebraic process calculi. J. Log. Algebr. Program. 73(1-2), 70–96 (2007)
20. Svensson, H., Fredlund, L.A., Earle, C.B.: A unified semantics for future Erlang. In Proc. of the 9th ACM SIGPLAN workshop on Erlang, pp. 23–32. ACM (2010)
21. Thomsen, M.K., Axelsen, H.B.: Interpretation and programming of the reversible functional language RFUN. In Proc. of IFL 2015. Springer (2016), *to appear*
22. Tiezzi, F., Yoshida, N.: Reversible session-based pi-calculus. J. Log. Algebr. Meth. Program. 84(5), 684–707 (2015)
23. Yokoyama, T.: Reversible computation and reversible programming languages. Electronic Notes in Theoretical Computer Science 253(6), 71–81 (2010), proc. of the Workshop on Reversible Computation (RC 2009)
24. Yokoyama, T., Axelsen, H.B., Glück, R.: Principles of a reversible programming language. In Proc. of the 5th Conference on Computing Frontiers, pp. 43–54. ACM (2008)
25. Yokoyama, T., Axelsen, H.B., Glück, R.: Fundamentals of reversible flowchart languages. Theor. Comput. Sci. 611, 87–115 (2016)